

Section II

Intel BTX Specification V 1.0 (assorted)

Balanced Technology Extended (BTX) System Design Guide

Version 1.0

2.3 Concurrent Design Engineering

The BTX form factor was concurrently developed and engineered by Intel® motherboard and system engineering groups. Motherboard component placement, therefore, takes into account not only the routing and electrical performance requirements, but also the heat transfer and airflow management required to ensure a proper operating environment.

Intel understands the migration from an existing to a new form factor standard is one that will create transition challenges, in addition to new opportunities for innovation. The new constraints imposed by the BTX form factor are ones that Intel anticipates will, in fact, lead to the use of the Desktop Personal Computer in more environments and usage models than is allowed by the existing form factors.

Intel has engineered BTX with the anticipation that all components will continue to evolve and that new components will need to be integrated. For instance, the motherboard core area – processor, chipset, and memory – has been increased to not only decrease component and routing density on today's platform but also for the anticipated growth in component power, I/O count, I/O density, and component physical size. By providing this additional board real estate, it is anticipated that standard Desktop systems will continue to use four-layer motherboard technology and cost effective component packaging technologies.

2.4 BTX Airflow Management Strategy

2.4.1 Provide the High Power Components with Low Temperature Air at High Velocity

The benefits of providing low temperature air at high velocity to the high power components include lower heatsink costs, improvement in acoustic performance, and the potential for improved component performance.

For a heatsink attached to a component with Heat Flux (Q) and a Package Case Temperature (T_{case}) specified, the heatsink resistance requirement is defined at the Design Ambient Temperature ($T_{ambient}$) as:

Equation 1: Thermal Characterization Parameter (Case-to-Ambient)

$$\psi_{ca} = [T_{case} - T_{ambient}] / Q$$

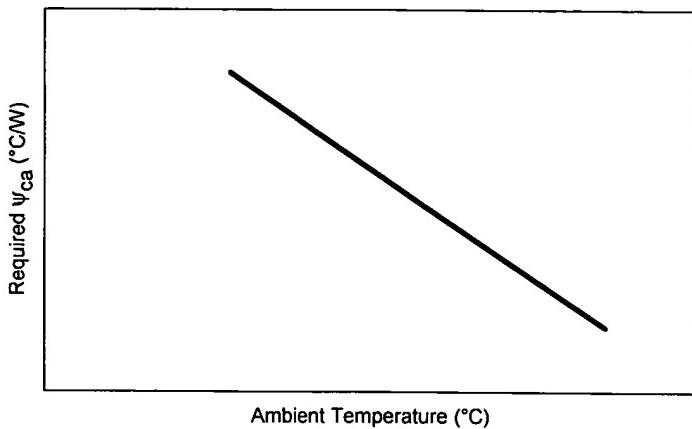
Similarly, the component may have a silicon junction temperature requirement instead of a case temperature requirement. In that case the heatsink resistance requirement is defined as:

Equation 2: Thermal Characterization Parameter (Junction-to-Ambient)

$$\psi_{ja} = [T_{junction} - T_{ambient}] / Q$$

As can be seen from the governing equation, for a given power into a heatsink and a given component temperature requirement (either T_{case} or $T_{junction}$), lower ambient temperature ($T_{ambient}$) increases the required heatsink resistance, as shown in Figure 1.

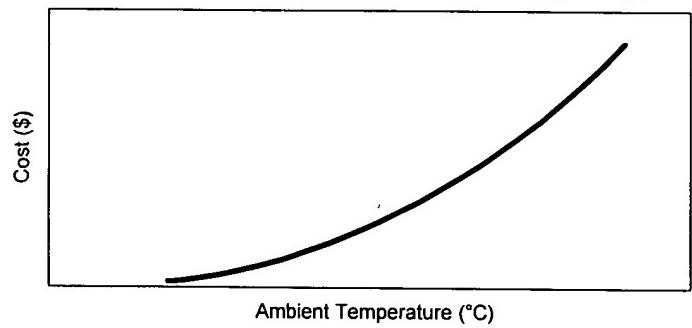
Figure 1: Requirement versus Ambient Temperature



OM16706

Higher resistance heatsinks are generally easier and cheaper to manufacture. The relationship between T_{ambient} and heatsink cost is shown in Figure 2 and illustrates why BTX was designed to provide lower temperature air to all high power Desktop system components.

Figure 2: Heatsink Cost versus Ambient Temperature



OM16707

Intel provides the required junction or case temperature specification for its processor and chipset components in component datasheets. Intel also provides guidance on designing and measuring the performance of heatsinks for these components in the component design requirements documents and thermal design guides.

Heatsink performance can be defined in terms of its conduction and convection performance. Conduction performance is primarily determined by the selection of heatsink materials (thermal conductivity, k) and design features such as its cross-sectional conduction area (A_c) and the length of the conduction path (L), as illustrated in Equation 3 by the governing equations for one-direction conduction.

Equation 3: Heat Conductive Performance

$$Q = k A_c \Delta T / \Delta L$$

The governing equations for convection will be used to illustrate the strong relationship between the velocity of the air that moves through the heatsink fin channels and the convection performance.

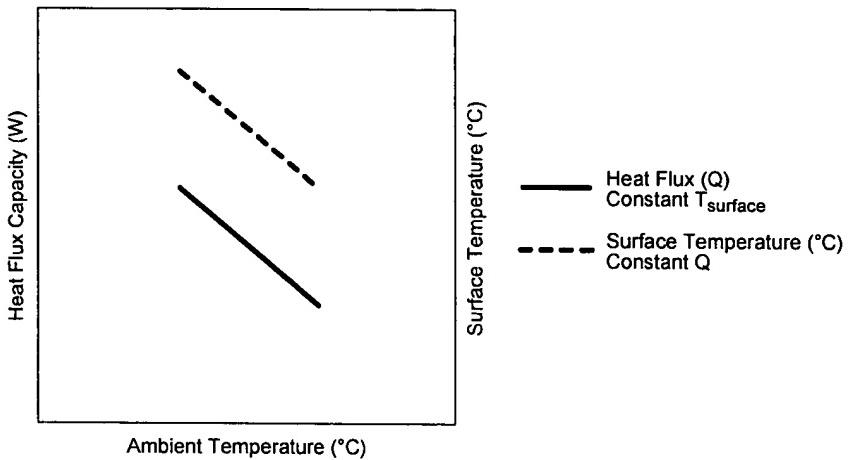
First, the amount of heat that can be dissipated by any convective surface is a function of the surface area from which it is being conducted (A), the temperature difference between the surface and the surrounding air (ΔT), and the convective coefficient (h).

Equation 4: Heat Convective Performance

$$Q = h A \Delta T = h A (T_{\text{surface}} - T_{\text{ambient}})$$

This equation illustrates the impact of lowering the ambient temperature. It is desirable to decrease T_{ambient} because it increases the amount of heat (Q) the surface can dissipate while maintaining a target surface temperature. Alternately, decreasing T_{ambient} can be used to lower the surface temperature for the targeted heat dissipation (Figure 3) while using the same thermal ingredient design.

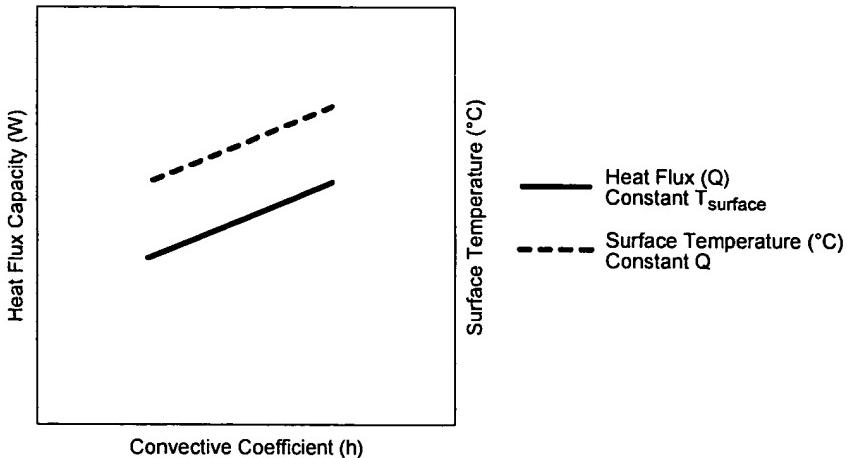
Figure 3: Effect of Decreasing Ambient Temperature



OM16708

Equation 4 also illustrates the impact of increasing the convective coefficient, h . It is desirable to increase h because it increases the amount of heat (Q) the surface can dissipate while maintaining a target surface temperature. Alternately, increasing h can be used to lower the surface temperature for the targeted heat dissipation (Figure 4).

Figure 4: Effect of Increasing Convective Coefficient, h



OM16709

The convective coefficient, h , is a function of the surface geometry, material conductivity (k), Reynolds Number (Re), and Prandtl Number (Pr):

Equation 5: Convective Coefficient

$$h = 0.664 k Re^{1/2} Pr^{1/3} / L$$

Note: The constant, 0.664, in Equation 5 is a function of the flow regime; this value increases when the flow transitions from the laminar to the turbulent regime. The flow regime in most areas of a Desktop system will be laminar, so the use of this equation is appropriate.

This equation is limited in application to the case where the air is flowing over a flat plate (for instance, over the surface of the motherboard, add-in card, or drive bay). For flow within a narrow channel, the boundary layer development will limit the improvement in convection coefficient due to increasing velocity.

For reference, the Prandtl Number is a function of the air viscosity (ν) and the thermal diffusivity of air (α):

Equation 6: Prandtl Number

$$Pr = \nu / \alpha$$

The Reynolds Number is a function of airflow velocity (V), the length of the surface (L), and air viscosity (ν). The Reynolds Number for flow over a horizontal surface (not the flow within a channel) can be calculated using the following equation:

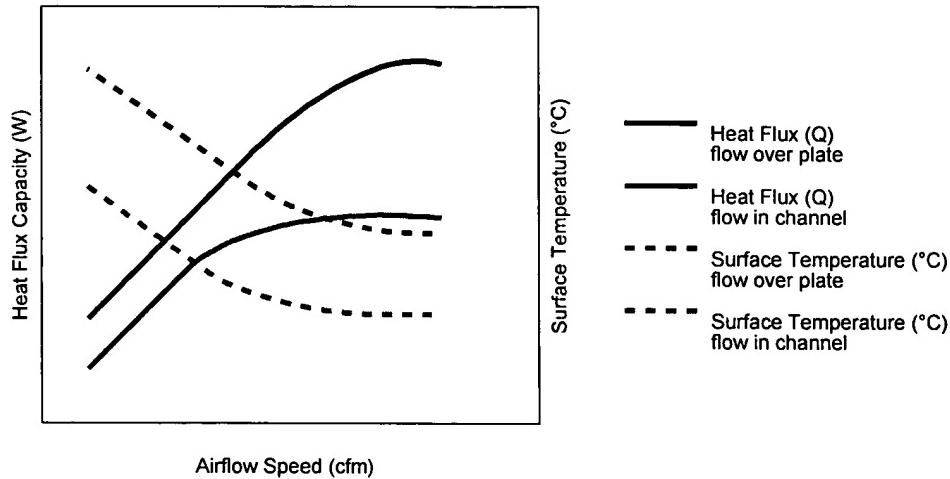
Equation 7: Reynolds Number

$$Re = V L / \nu$$

From Equation 5 and Equation 7, it can be demonstrated that the convective coefficient, h , is proportional to airflow velocity. Therefore, the convective heat transfer capability of a surface increases with an increase in velocity. Generally, it can be stated that heatsink resistance, ψ , decreases with increasing velocity.

At a target surface temperature, the heat power that can be dissipated increases with increasing airflow velocity through a heatsink. Alternately, at a target heat power, the surface temperature decreases with increasing airflow speed.

Figure 5: Effect of Increasing Airflow Velocity



OM16710

In summary, the performance of every convective surface is a function of the ambient airflow temperature and velocity. Lower airflow temperature increases ΔT and higher airflow velocity increases the convection coefficient. The BTX has been engineered to reduce the ambient temperature and increase the airflow velocity in the region of all high power components.

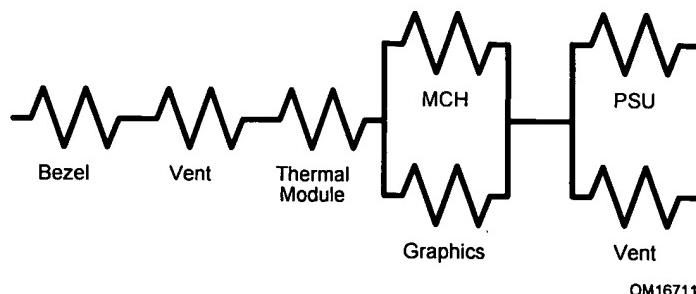
2.4.2 Minimize the Total System Impedance

Impedance is the resistance to movement that a moving fluid encounters. Generally, low impedance is more desirable than high impedance because it allows the air movers (e.g., fans) to provide higher airflow at each operating speed.

In Desktop systems, there are impedance sources that are more preferable than others. For instance, when air is flowing through a channel, the walls of that channel represent impedance to the airflow, but if the fins of a heatsink create that channel then the air is also working to remove heat. In other words, impedance introduced to remove heat or improve heat transfer performance is “good” impedance. Examples of undesirable impedance include airflow turning or forcing air through narrow channels and expansions or constrictions that do not remove heat from components.

A network resistance diagram should be constructed for every Desktop system design, illustrating the parallel and series airflow impedances from the airflow entrance to its exit. The impedance (resistance) characteristics of each subsystem can be estimated using hand calculations, or loss coefficient estimates from fluid flow handbooks, measured in a wind tunnel, or predicted using fluid dynamic numerical tools. A typical, simple resistance network diagram for a Desktop system is shown in Figure 6.

Figure 6: Typical Network Resistance Diagram



The total resistance to airflow in a Desktop system determines that amount of airflow that the various fans can deliver. Lower impedance systems allow more airflow to pass through them and, as seen from the previous section, this will improve the system's heat transfer performance.

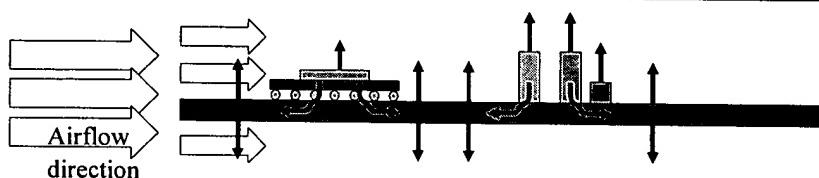
Acoustic performance is strongly correlated to the system impedance. That is, for any required system airflow, higher system impedance will require higher fan speeds and, therefore, create more acoustic noise. The relationship between impedance and the number and operating speeds of the fans is described in Section 1.

2.4.3 Provide Above-board and Under-board Airflow

There are several sources of heat within and directly attached to a motherboard. Metal traces and planes in the motherboard generate heat during the conduction of electrical current. Surface Mount components directly attached to the motherboard also generate and transfer heat into the motherboard. Motherboard reliability is compromised when its temperature exceeds operating specifications; therefore, removing heat from the motherboard is an important part of a Desktop system's thermal design.

By providing airflow above and below the motherboard, two convective heat transfer paths from the motherboard are established (effectively doubling the available convective heat transfer surface area relative to that in an ATX system), as shown in Figure 7. The improvement in total heat transfer capability from the motherboard allows higher power components to be attached to the motherboard without compromising motherboard reliability.

Figure 7: Two Heat Transfer Paths from Motherboard



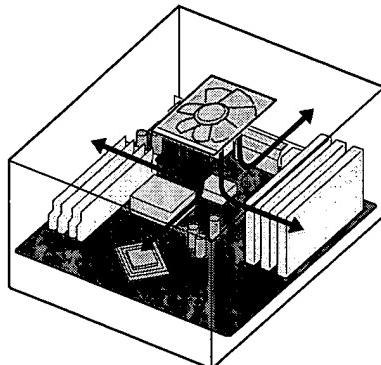
2.4.4 System Airflow Direction

The three important airflow management objectives identified in Sections 2.4.1, 2.4.2, and 2.4.3 were used by Intel system and motherboard engineers in the placement of high power components in a single airflow stream generated and managed by a minimum number of fans. With this airflow management and component placement strategy outlined, the next challenge was to determine the most appropriate direction for that airflow.

Creating a single airflow stream direction implies that the Desktop system must be designed with the inlet and exhaust at opposite sides of the chassis.

The path from the chassis pan to top cover was not considered feasible because airflow would impinge on the motherboard, as shown in Figure 8. The associated turning impedance is not “good” impedance, so this flow direction was not selected.

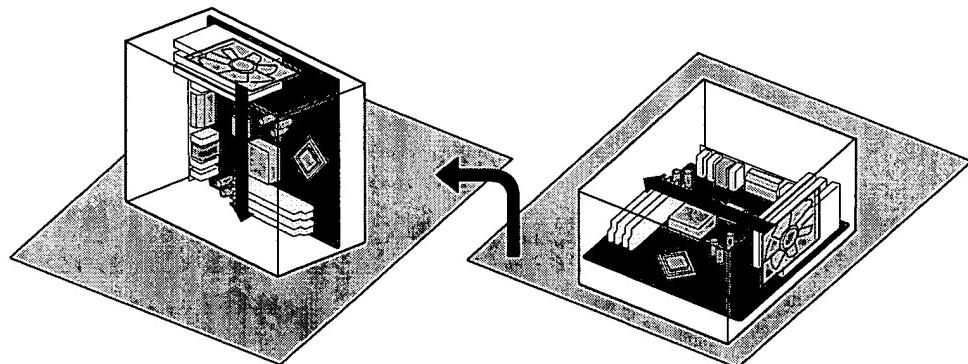
Figure 8: Top-to-Bottom Airflow Illustration – Impinging Flow at Motherboard



OM16713

Since it is common in Desktop system design to have a single chassis design operate in desktop and tower orientations, inlet and exhaust from side-to-side or top-to-bottom would have the inlet or exhaust blocked by the surface on which the chassis rests (Figure 9); therefore, these airflow directions were not selected.

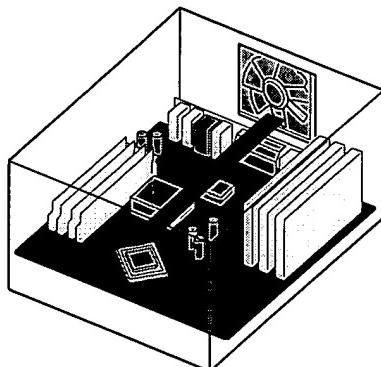
Figure 9: Side-to-Side Airflow Illustration – Exit Flow Blocked When System Is Rotated



OM16714

Finally, exhausting airflow out of the front panel of the chassis directs warm air exhausted from the system toward the user, as shown in Figure 10.

Figure 10: Back-to-Front Airflow Illustration – Exit Airflow Directed at User



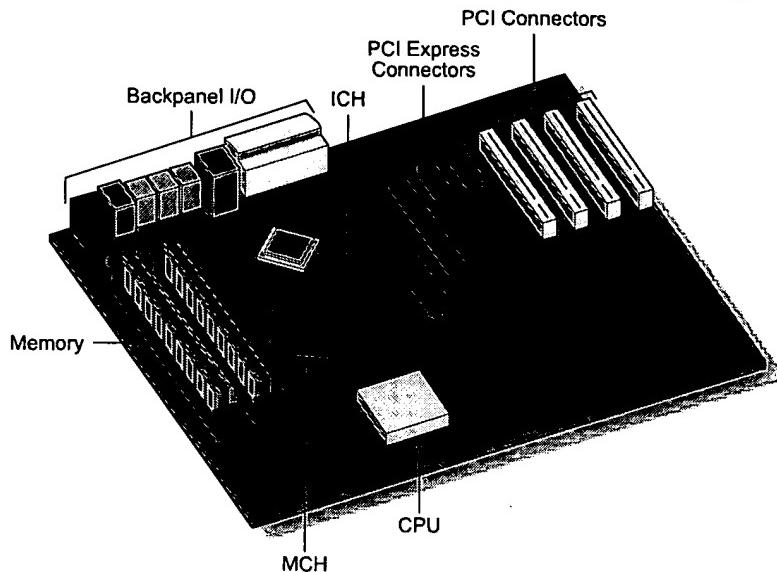
OM16715

Therefore, Intel designed BTX to use front-to-back airflow with the high power components aligned so that a single airflow stream services them all. Placing a fan at the front of the system creates unique acoustic concerns because the fan is now a noise source that is directly in front of the system user. Acoustic management strategies for a noise source at the front of the system are discussed in the Acoustic Engineering section of this design guide.

2.4.4.1 CPU Location

The highest power component with the most stringent temperature specifications is the Central Processing Unit (CPU), so it is located in the front of the motherboard, near the chassis front panel (typical component locations on a BTX motherboard are illustrated in Figure 11). In this position, the CPU will receive the lowest ambient temperature air moving at its highest speed, as described in the front-to-back airflow management strategy outlined above.

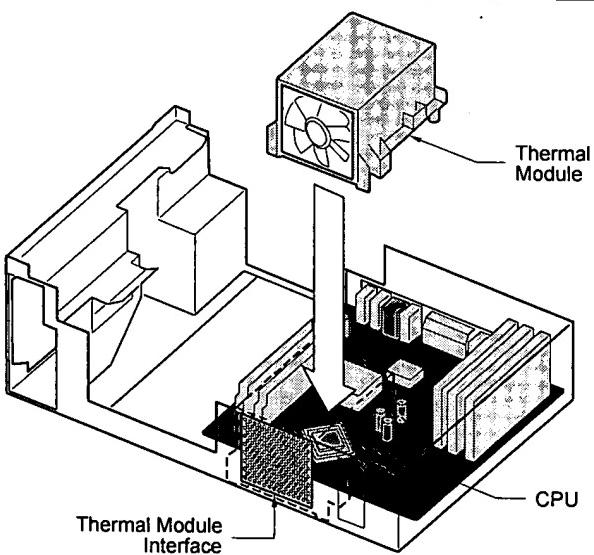
Figure 11: Typical Component Locations on a BTX Motherboard



OM16716

The fan in the BTX Thermal Module provides the airflow for the CPU heatsink and the remainder of the system. Figure 12 illustrates the position of the Thermal Module relative to the CPU.

Figure 12: Thermal Module Location Relative to the CPU



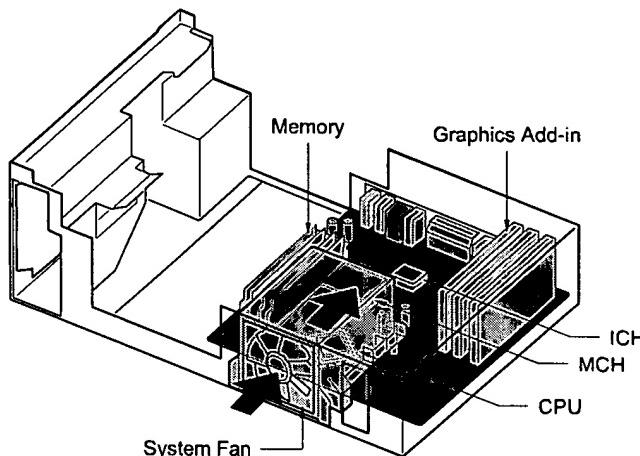
OM16717

The *BTX Interface Specification* requires that the Thermal Module Interface be designed in such a way as to eliminate the potential for air that has entered the system to re-enter the Thermal Module inlet. Eliminating recirculation ensures that the CPU heatsink receives the lowest temperature air, which comes from outside the chassis. It also ensures that the temperature of Thermal Module exit airflow (which

provides the airflow for the remainder of the system) is not increased by recirculation. Figure 12 also illustrates a Thermal Module Interface provided by a vented portion of the chassis sheet metal front panel.

The Thermal Module supplier is responsible for providing a fan, in addition to the CPU heatsink and duct assembly. The Thermal Module exhaust air – which exits the CPU heatsink and duct assembly – provides the airflow for the remainder of the Desktop system components. Figure 13 illustrates the primary airflow path generated by and exiting from the Thermal Module fan.

Figure 13: Thermal Module Airflow Path



OM16718

With a properly selected Thermal Module, the exhaust airflow will provide low temperature air at high velocity to the remainder of the system components. The importance of selecting a Thermal Module with an appropriate effective fan curve is discussed in Section 2.4.5.1.

The location of the CPU at the front of the system is beneficial not only to the CPU heatsink performance, but also to the CPU socket performance. The CPU socket generates heat when it is delivering the operating current required by the CPU. In BTX, the temperature of the socket and motherboard are reduced because there is low temperature, high velocity airflow directed above the motherboard at the socket and because the bottom of the motherboard is a convective surface due to the under-board airflow (refer to Figure 7).

Comparison to ATX: ATX CPU thermal solutions typically have a dedicated fan directing airflow through a heatsink. This impinging flow configuration forces airflow to turn at the motherboard, which introduces substantial undesirable impedance to the CPU heatsink fan. This increase in impedance decreases the velocity of the air through and exiting the ATX CPU heatsink. The exhaust airflow path is constrained by the location of memory and the first expansion slot, forcing the exhaust air to recirculate into the CPU heatsink fan. The increase in CPU heatsink inlet temperature decreases its heat transfer capability. The increase in exhaust airflow temperature also decreases the heat transfer efficiency around the socket and motherboard. A BTX Thermal Module will not have turning impedance or recirculation; therefore, the CPU heatsink and socket-motherboard heat transfer efficiency will be higher.

Refer to Section 2.4.5.1 for more information on the Thermal Module.

2.4.4.2 CPU Voltage Regulation Location

CPU Voltage Regulation component performance is typically constrained by the Voltage Regulation component or motherboard temperature specifications. The electrical efficiency losses that result from the management of the CPU operating current create heat in the Voltage Regulation components and in the motherboard traces and power planes.

CPU Voltage Regulation is typically designed using multiple phases; each phase is comprised of Voltage Regulation components (drivers, capacitors, inductors, etc.) and is responsible for managing a portion of the CPU operating current. When the Voltage Regulation component or motherboard temperature specifications are exceeded because the heat associated with the phase's operating current, the following design options are typically considered:

- a. Attach heatsinks to the Voltage Regulation components to improve their convective heat transfer performance and thereby lower their operating temperature.
- b. Provide more airflow to Voltage Regulation components to improve their convective heat transfer performance and thereby lower their operating temperature.
- c. Add a phase to the Voltage Regulation design to reduce the operating current per phase and thereby reduce the amount of heat generated per phase.
- d. Redistribute the current going to each phase such that it is not equally distributed.

The BTX Thermal Module provides above and below motherboard airflow at low temperature. This airflow pattern creates the heat transfer illustrated in Figure 6. CPU Voltage Regulation components are typically placed near the CPU socket. The greatest heat transfer benefit will be realized if these components are placed in front of the CPU socket, near the front of the BTX motherboard because they will receive low temperature air directly from the Thermal Module fan.

Comparison to ATX: The exhaust airflow from an ATX CPU thermal solution typically provides the airflow for cooling CPU Voltage Regulation and the motherboard. As discussed in Section 2.4.4.1, this impinging flow configuration decreases the velocity and increases the temperature of the heatsink's exhaust airflow. This low velocity, high temperature airflow on only one side of the motherboard does not create an efficient heat transfer path for CPU Voltage Regulation or the motherboard. A BTX Thermal Module provides the same high velocity, low temperature airflow to the CPU Voltage Regulation, socket, and motherboard as it does to the CPU heatsink. It provides this airflow above and below the motherboard, creating a very efficient heat transfer path.

For more information on engineering CPU Voltage Regulation airflow, refer to Sections 2.4.5.2 and 2.4.5.3.

2.4.4.3 Graphics Location

Integrated graphics is often included on Intel® Memory Controller Hub (MCH) chipsets. Refer to Section 2.4.6 for a description of the MCH airflow environment.

A graphics add-in card will typically be located in the first add-in card slot position. In BTX, the graphics add-in card slot position was shifted to the right side of the core so that the powered components of the card would be exposed to the exhaust airflow of the BTX Thermal Module. The card is aligned with the exhaust airflow direction so that it does not create unnecessary system impedance.

Whether the card is placed directly in the first slot (perpendicular to the motherboard) or on a riser (parallel to the motherboard), the powered components of the add-in card – the Graphics Processing Unit, Voltage Regulation, and memory – are exposed to the Thermal Module exhaust airflow.

Comparison to ATX: In ATX, the powered components on a graphics add-in card are on the side opposite the CPU, and are facing into a very narrow channel created by the adjacent card. There typically is very little system-generated airflow on this side of the card and it is not uncommon for a performance ATX graphics add-in card to include a thermal solution with a heatsink and dedicated fan. The high velocity, low temperature exhaust airflow provided by the BTX Thermal Module moves past the powered side of a graphics add-in card, significantly improving the heat transfer performance of the card's thermal solution. The improvement allows lower graphics card heatsink cost, improved operating temperature, or higher power.

For more information on engineering graphics add-in card airflow, refer to Section 2.4.9.

2.4.4.4 Memory Controller Hub (MCH) Location

In BTX, the MCH is located immediately behind the CPU on the motherboard. In this position it will receive the high velocity, low temperature Thermal Module exhaust airflow. That airflow direction will be consistent in every BTX system so the MCH heatsink fins can be continuous instead of crosscut.

In addition, a conductive heat transfer path typically exists between the MCH and motherboard - the MCH solder joints conduct heat from the MCH silicon and package into the motherboard. This heat flow into the board is eventually transferred by convection from the topside of the board but generally increases the board temperature. The above-board and under-board airflow provided in a BTX system provides two convective heat transfer paths for this heat to leave the motherboard.

Comparison to ATX: The exhaust airflow from an ATX CPU thermal solution typically provides a portion of the airflow for the MCH heatsink. This impinging flow configuration decreases the velocity and increases the temperature of the CPU heatsink's exhaust airflow. An ATX system fan may also provide a portion of the airflow for the MCH heatsink, but the direction of that airflow is not consistent in all ATX system designs; therefore, an MCH heatsink is typically crosscut. The high velocity, low temperature exhaust airflow provided by the BTX Thermal Module directly to the MCH heatsink significantly improves its heat transfer performance. And the below board airflow offers another MCH heat transfer path. These improvements in heat transfer allow lower cost MCH heatsinks, improved operating temperature, or higher power.

For more information on engineering MCH airflow, refer to Section 2.4.6.

2.4.4.5 Memory Location

Much the same as the graphics add-in card slot, memory is aligned with the Thermal Module airflow to minimize the impedance it introduces. Thermal Module exhaust airflow will be drawn into and through the memory channels by the power supply fan, if the PSU is located near the memory fixed motherboard edge. In addition to Thermal Module exhaust airflow, memory airflow can be augmented by the airflow that exhausts below the CPU heatsink from the CPU Voltage Regulation. A BTX Thermal Module should be designed such that the below-heatsink airflow path is not fully ducted from the heatsink base to the motherboard. The below-heatsink airflow that is directed toward memory by the socket is an important system airflow path.

Comparison to ATX: The exhaust airflow from an ATX CPU thermal solution typically provides a portion of the airflow for memory. An ATX system fan may also provide a portion of the memory airflow but the direction of that airflow is not consistent in all ATX system designs. BTX Thermal Module and CPU Voltage Regulation exhaust airflow will provide consistent inlet flow to the memory airflow; however, managing memory airflow bypass and exit is important to ensuring an appropriate thermal environment.

For more information on engineering memory airflow, refer to Section 2.4.9.

2.4.5 Subsystem and Component Airflow Management Strategies

2.4.5.1 Thermal Module Role

The Thermal Module fan is the primary air mover in a BTX system. Not only does the Thermal Module fan supply airflow directly to the CPU heatsink, CPU socket, and CPU Voltage Regulation components, its exhaust airflow is used as a medium to provide the heat transfer for the remaining subsystems and components.

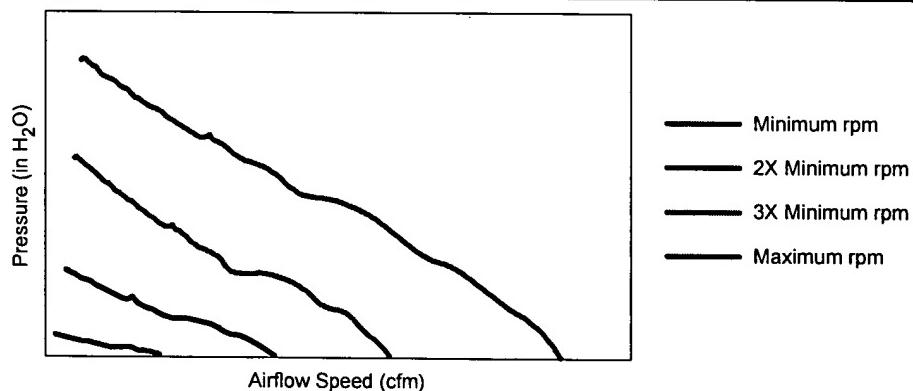
In fact, a typical BTX chassis will not ship with any fans installed since the system airflow will be provided by the Thermal Module and PSU fans. The Thermal Module fan will be provided by the Thermal Module supplier and the PSU fan will be provided by the PSU supplier.

2.4.5.1.1 Thermal Module Effective Fan Curve

The Thermal Module exit airflow characteristics are described by its effective Thermal Module fan curve, which is a function of the stand-alone fan curve, the CPU heatsink impedance, CPU Voltage Regulation impedance, CPU socket impedance, and any near-field impedance effects.

Measuring the volume of airflow that a fan can generate against increasing impedance generates a stand-alone fan curve. This measurement is typically conducted in a wind tunnel. The stand-alone fan curve varies with its operating speed, as shown in Figure 14. All fan curves reflect the following performance characteristic: at a particular fan operating speed, as the impedance in a fan's airflow path increases, the amount of airflow that the fan can generate decreases. A stand-alone fan curve at the maximum operating speed is typically in the fan performance data sheet provided by the fan supplier.

Figure 14: Fan Performance Curve versus Fan Speed (RPM)



OM16719

Fans typically include a closed-loop feedback design that monitors and maintains the requested speed. The impact of the closed-loop feedback circuit on the fan curve is best described using the following sequence of events:

- a. A fan converts electrical power into the momentum of the air it is attempting to move. The more difficult the air is to move, the more electrical power is required to impart that momentum.
- b. Impedance at the exit of a fan increases the difficulty of imparting momentum to the air, because the increase in exit impedance increases the pressure in the fan. The increased pressure in the fan makes it more difficult for the impeller to move through and impart momentum to the air.

- c. If the electrical power input to a fan is fixed, then increasing the exit impedance causes the fan to slow down. That is, since it finds the air more difficult to move, the rotation of the fan is slowed.
- d. The fan's closed loop feedback circuit has a tachometer that monitors fan speed. When higher impedance slows a fan down, the circuit will attempt to return the fan speed to the requested speed. It attempts to do this by increasing the electrical current supplied to the fan motor windings.
- e. If the fan feedback circuit request for more electrical current can be satisfied, the fan's operating speed will be maintained, even as impedance increases. In an operating regime where the higher current request can be met, the fan will operate at constant speed. The fan curve in this regime is termed an Iso-RPM fan curve.

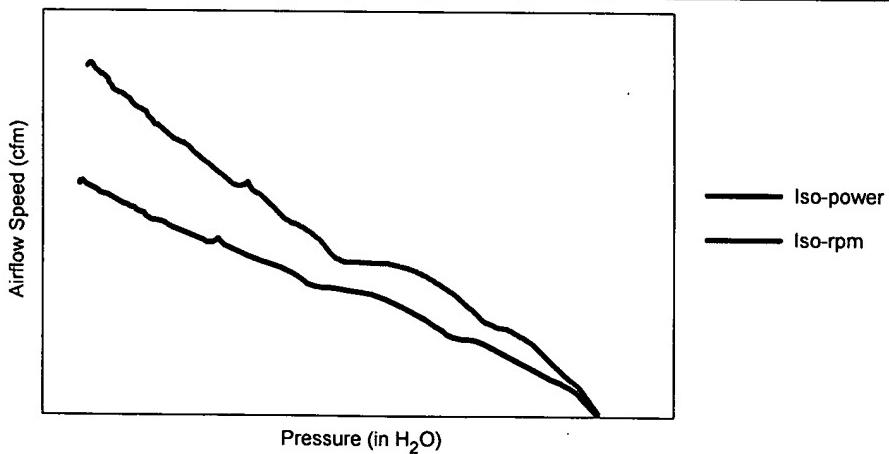
Note: Even if the fan is able to maintain its operating speed against increasing impedance, the amount of airflow that can be generated decreases with increasing impedance. That is, the impedance affects the volume rate of air that the fan can move even if the fan speed can be maintained. This is the performance characteristic reflected in the fan performance curve (see Figure 14).

- f. The current required to operate a fan increases as the requested fan speed increases and as the impedance increases.
- g. Therefore, at higher operating speed and higher impedance, the current requested by the fan motor may eventually exceed the current that can be provided.

For instance, when the fan is operating at its maximum speed against low impedance it will draw electrical current that may be near the current limit of the electrical circuit. As the impedance increases, the fan will request more electrical current in an attempt to maintain the operating speed. However, once the current limit of the circuit is reached, the fan will no longer be able to maintain the operating speed. If the impedance continues to increase, the fan performance curve will reflect the decrease in airflow volume rate (CFM – cubic feet per minute) from both the impedance increase and the reduction in fan speed.

Figure 15 shows the impact of the reduction in fan speed from this constrained electrical current operating condition. One curve shows the fan curve if electrical current were unconstrained; the other shows the reduction in CFM after the current limit is reached. The fan curve in this current limited operating condition is called an iso-power fan curve.

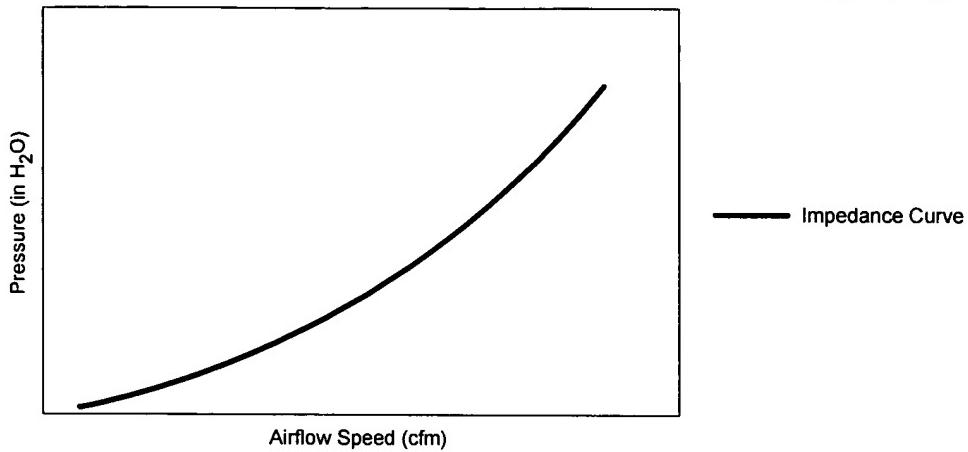
Figure 15: Iso-RPM versus Iso-Power Fan Performance Curve



OM16720

The impedance of any particular airflow path changes with the amount of airflow being forced into it. This characteristic of the airflow path is called its impedance curve. An impedance curve can be determined through fluid dynamic numerical modeling applications or wind tunnel measurement. A typical impedance curve is illustrated in Figure 16.

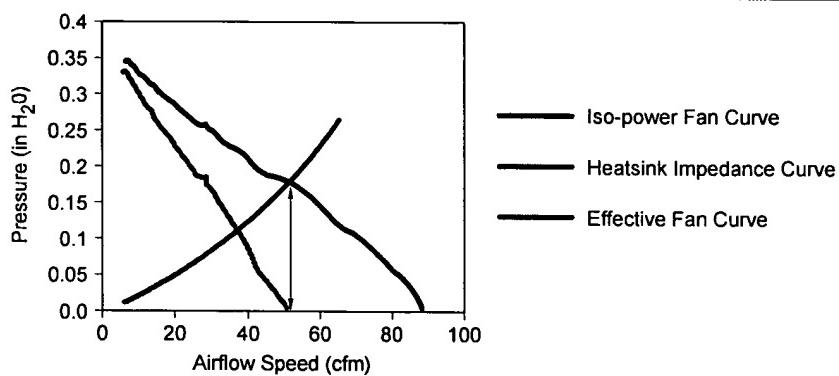
Figure 16: Typical Impedance Curve



OM16721

The Thermal Module effective fan curve is an indication of the airflow that the Thermal Module can provide a system after the airflow has made its way through the CPU heatsink, Voltage Regulation, and socket. This effective fan curve reflects the impact that the CPU heatsink and Thermal Module duct have on the stand-alone fan curve. Theoretically, if the impedance of the CPU heatsink and duct were measured in a wind tunnel then subtracted from the measured or supplier-provided stand-alone fan curve, the result would be the Thermal Module effective fan curve. This is illustrated in Figure 17.

Figure 17: Theoretical Thermal Module Effective Fan Curve Calculated From Stand-alone Fan Curve and Heatsink Impedance Curve



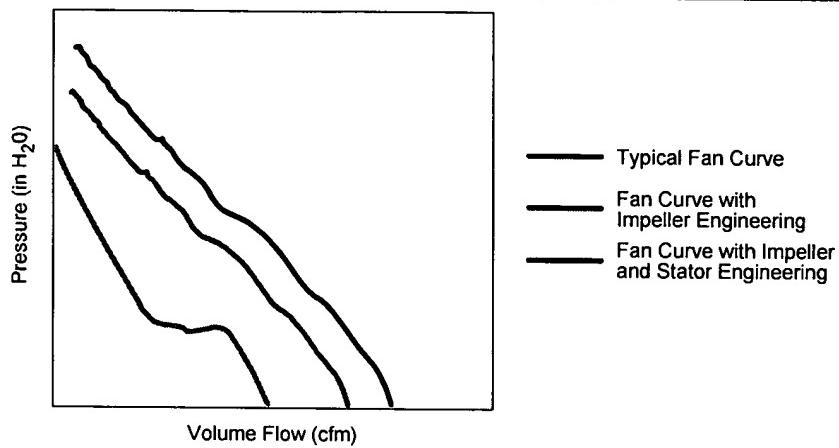
OM16722

However, the near-field impact that the CPU heatsink has on the stand-alone fan curve is difficult to accurately predict from the stand-alone fan curve and the heatsink impedance curve. This is because a wind tunnel can measure only the axial flow characteristics; whereas a rotating radial fan creates airflow

with axial and swirl components. When the swirl component of the fan airflow enters the heatsink, it creates additional impedance that is not present in the axial wind tunnel measurement of the fan curve or the heatsink impedance. It is best to measure the Thermal Module effective fan curve by inserting the entire Thermal Module (fan, heatsink, and duct on a test board) into the wind tunnel.

Increasing the axial component and minimizing the swirl component of airflow that exits the Thermal Module fan can minimize the impact that the CPU heatsink has on the effective fan curve. The Thermal Module supplier can select or engineer the fan impeller to increase the axial component of airflow, and select or engineer a stator that converts a portion of the swirl component to axial. The improvement in the stand-alone fan curve through this type of fluid dynamic engineering is illustrated in Figure 18.

Figure 18: Fan Curve Impact from Impeller and Stator Fluid Dynamic Engineering

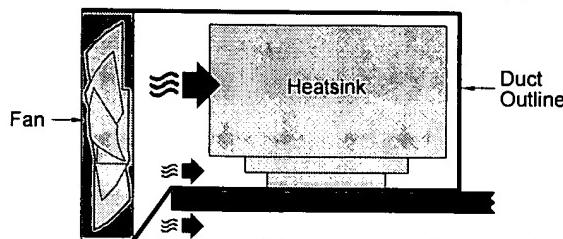


OM16723

2.4.5.1.2 Primary Airflow Paths Within Thermal Module

The Thermal Module is responsible for partitioning the airflow to three distinct airflow paths: through the ducted CPU heatsink, above the motherboard but below the CPU heatsink base, and below the motherboard (see Figure 19).

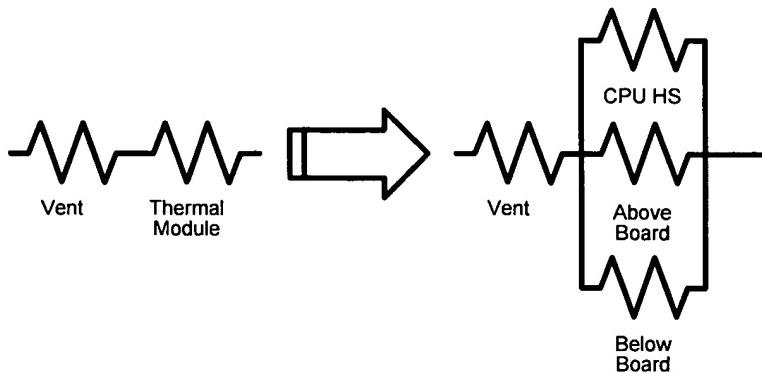
Figure 19: Thermal Module Airflow Partitioning



OM16724

The network resistance diagram in Figure 6 can be refined to show the partitioning of these airflow paths in the Thermal Module, as illustrated in Figure 20. This is an excellent method for identifying and managing airflow partitioning and the specific airflow volume through each path.

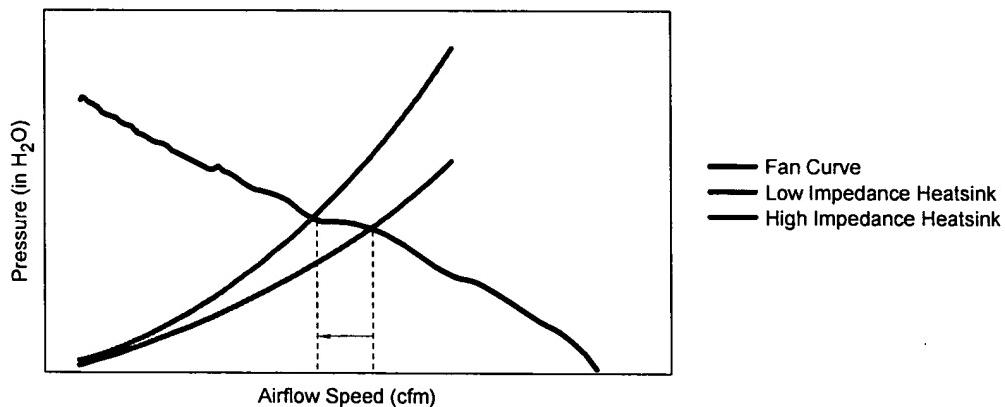
Figure 20: Thermal Module Network Resistance Diagram



OM16725

The airflow through the CPU heatsink is determined by the interaction between the Thermal Module fan stand-alone fan curve and impedance characteristic of the selected CPU heatsink design (as illustrated in Figure 17). Improvements in CPU heatsink heat transfer performance often involve adding material to the heatsink. Improvements in conduction can be accomplished through increases in the conduction path cross-sectional area, and improvements in convection can be accomplished through increases in surface area and the addition of fins. Both of the engineering approaches typically increase the heatsink impedance because they decrease the channel area available for airflow. The reduction in airflow from an increase in heatsink impedance is illustrated in Figure 21 by the superposition of the heatsink impedance curves and the stand-alone fan curve.

Figure 21: Impact of Heatsink Impedance on Heatsink Airflow



OM16726

The reduction in airflow not only impacts the CPU heatsink heat transfer performance by reducing its convection coefficient (see Equation 4), but it also impacts the Thermal Module effective fan curve (see Figure 17). Recall that the Thermal Module Effective fan curve is an indication of the airflow and pressure capability available for the remainder of the system. It is also important to note that the CPU heatsink impedance also influences the airflow that will go through the remaining paths; a higher impedance heatsink may force more airflow into the above and below motherboard airflow paths.

A Thermal Module supplier will need to carefully select or engineer the fan and heatsink characteristics to achieve the required CPU heatsink heat transfer performance (ψ_{ca}) and an appropriate effective fan curve. Returning to the network resistance diagram in Figure 20 and using the governing equations for airflow as a function of impedance (Equation 8), it is obvious that an increase in the impedance of one of the paths in parallel may increase the airflow to the remaining paths.

Equation 8: Airflow as a Function of Impedance

$$W = [2 \cdot dP / ((L_1 + L_2) \cdot \rho)]^{1/2} \quad \text{Flow through channels in series}$$

$$W = [2 \cdot dP / (\rho \cdot (1/\sqrt{L_1} + 1/\sqrt{L_2})^2)]^{1/2} \quad \text{Flow through channels in parallel}$$

Where:

W is the volumetric flow rate

dP is the pressure loss through two channels in series

L₁ and L₂ are the length of the two channels

ρ is the density of the fluid (air)

Additional information on selecting or engineering above and below motherboard airflow impedance is discussed in Sections 2.4.5.2 and 2.4.5.3.

A system integrator will need to select the Thermal Module based on the required CPU heatsink performance and the required effective fan curve.

2.4.5.2 Above Motherboard Airflow Path

The airflow for this path enters the channel created by the base of the CPU heatsink and the motherboard. As stated previously, the airflow speed will be a function of the Thermal Module stand-alone fan curve and the impedance of the three primary airflow paths: above and below the motherboard and through the ducted CPU heatsink.

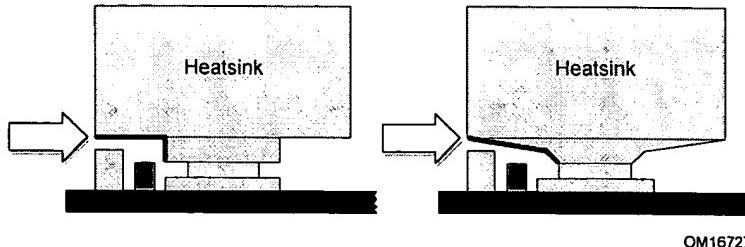
The impedance of the above board airflow path is typically governed by the following physical characteristics:

- a. The dimension of the channel created by the CPU heatsink base and motherboard.

The impedance of the channel decreases as the channel height increases. Since the base of the CPU heatsink typically rests on the top of the CPU package, there are few design options to change this channel height.

The shape of the bottom surface of the CPU heatsink base will affect the impedance of this channel. For instance, in Figure 22, the heatsink base shown on the right will have slightly lower channel impedance than the one on the left due to the relative reduction in turning loss.

Figure 22: Heatsink Base Impact on Above Motherboard Channel Impedance

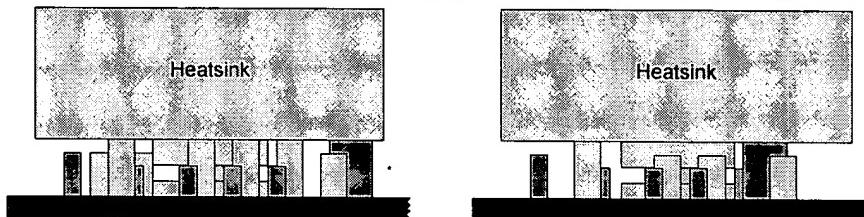


OM16727

- b. The size, quantity, and position of motherboard components near the CPU socket.

The impedance of the channel increases as the total frontal area of these components increases. The frontal area will increase as components are added, as those components become wider and taller, and when the components are closely spaced. For instance, in Figure 23 the Voltage Regulation configuration on the right will have slightly lower channel impedance than the one on the left (in this illustration, the airflow enters the page) because it has more open area available for the airflow that moves through it.

Figure 23: Voltage Regulation Component Impact on Motherboard Channel Impedance



OM16728

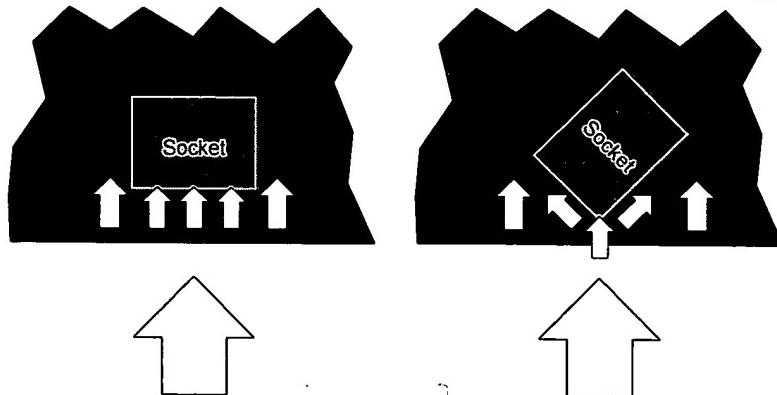
Since low profile Voltage Regulation components are typically more expensive, the system integrator will need balance the cost of the selection or specification of Voltage Regulation components against the airflow impedance that they will introduce.

- c. The size and orientation of the CPU socket.

A larger CPU socket will increase the channel impedance; however, the orientation of the socket has a more significant influence. If the socket edge is perpendicular to the airflow direction, its impedance will be at its highest possible value. As the socket is rotated such that its corner is presented to the airflow, its impedance reduces. For instance, in Figure 24 the impedance of the socket on the right will be lower than the one on the left due to the relative reduction in turning loss.

BEST AVAILABLE COPY

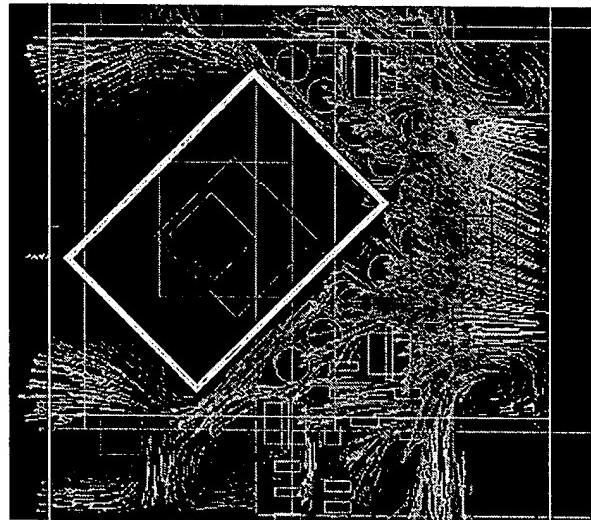
Figure 24: Socket Orientation Impact on Impedance



OM16729

Also, if the socket body is rotated, it acts as an effective airflow diverter. This allows the airflow to slipstream around the socket body and be efficiently directed to the remaining portions of the system. Figure 25 is an image extracted from a fluid dynamic numerical model of the airflow near a rotated socket body. In this illustration, the airflow comes from the Thermal Module fan that is mounted to the right. Note how the airflow efficiently begins to move around the socket body early in the airflow stream.

Figure 25: Rotated Socket Acting as a Flow Diverter



OM17114

The above motherboard airflow that goes around the socket and through the CPU Voltage Regulation components eventually provides inlet airflow to other portions of the system. Memory, located to the left of the CPU socket (in Figure 25, memory is located toward the bottom left of the illustration), receives a portion of the above motherboard airflow from the Voltage Regulation and socket area. This memory airflow path complements the Thermal Module exhaust airflow that exits the CPU heatsink.

2.4.5.3 Below Motherboard Airflow Path

The Thermal Module provides airflow below a BTX motherboard because the Thermal Module fan is located in front of, and its exit face extends below, the motherboard (see Figure 19). This airflow allows the bottom side of the motherboard to become a convective heat transfer surface. The motherboard is, of course, a complex electrical circuit composed of many conduits for managing electrical current (e.g., traces, planes, vias, and pads). As these motherboard features manage electrical current, that current generates heat power and, unless that heat is completely dissipated into the surrounding environment, the temperature of the motherboard will rise. When the motherboard temperature at any given location exceeds the operating temperature specification, the motherboard is subject to performance degradation. Since the above and below motherboard airflow in BTX allows heat to be effectively dissipated through its top and bottom surfaces, a BTX motherboard can manage higher current in critical locations. Especially in the region of the CPU socket and Voltage Regulation components, the bottom side heat transfer characteristic significantly increases the amount of heat, and thereby the amount of electrical current, that can be managed.

Surface mounted motherboard components can conduct a portion of the heat they generate into the motherboard. When the convective heat transfer capability of the motherboard is increased by above and below motherboard airflow, more heat from these surface mount components will conduct into the motherboard, effectively lowering the temperature of those components. This is especially important for CPU Voltage Regulation components. Components mounted to the bottom side of the motherboard will see the direct benefit of below motherboard airflow, since that airflow creates convective heat transfer directly from the components, in addition to removing heat from those components that is conducted into the motherboard. Even for components mounted on the topside of the motherboard, the below motherboard airflow is of significant benefit.

The below motherboard airflow is also a significant benefit for the Memory Controller Hub (MCH) and Input-Output Controller Hub (ICH) chipsets. Again, these components will conduct a portion of their heat into the motherboard and the convective heat transfer on the motherboard bottom side created by the below motherboard airflow effectively reduces the temperature of these components. More detail on MCH heat transfer is provided in Section 2.4.6 and on ICH heat transfer in Section 2.4.7.

The impedance of the below motherboard airflow channel will increase when the channel area is reduced. As described in Chapter 4, a Structural Retention Module is required in BTX system designs. Since this component occupies a portion of the space between the chassis pan and motherboard, it increases the below motherboard impedance and reduces the airflow. Components mounted on the bottom side of the motherboard will also increase the impedance.

Additional guidance on the thermal performance improvement from below motherboard airflow may be provided in subsequent revisions of this guide.

A system designer will need to pay particular attention to the way in which the Thermal Module airflow is partitioned. While the above and below motherboard airflow paths are important for managing CPU Voltage Regulation, CPU socket, and, to some extent, chipset temperatures, the CPU heatsink and the Thermal Module effective fan curve are typically more important. If the impedance of the above and below motherboard channels is particularly low, then more of the Thermal Module fan's airflow will go through these channels – meaning that less of that airflow will go through the CPU heatsink. Not only does this reduce the CPU heatsink heat transfer performance but it can also impact the heat transfer characteristics for the remaining portions of the system, since the Thermal Module effective fan curve essentially determines these important total system flow and flow distribution characteristics.

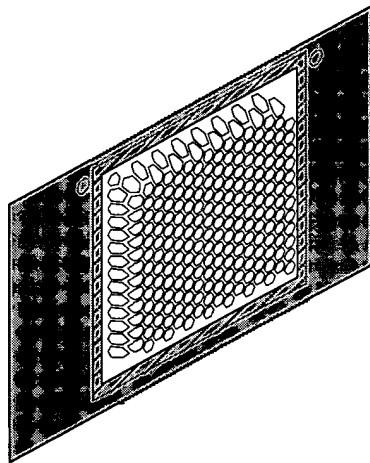
It may be necessary for the Thermal Module supplier to introduce impedance to the above and below motherboard channels to ensure that the appropriate airflow is directed to the CPU heatsink. System

integrators should select a Thermal Module based on its CPU heatsink ψ_{ca} and effective fan curve, but they should also understand the impact of Thermal Module design options on the above and below motherboard airflow.

2.4.5.4 Airflow into the Thermal Module

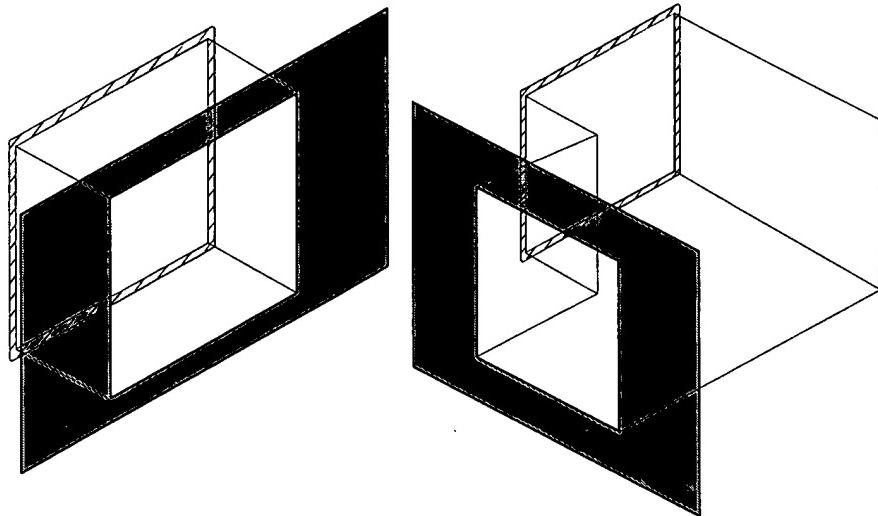
The *BTX Interface Specification* defines a requirement for a Thermal Module Interface. Chassis suppliers must always provide a Thermal Module Interface that is sealed and meets the open area, positional, and size requirements defined in the *BTX Interface Specification*. A BTX Form Factor compliant Thermal Module Interface may be provided through the use of a vented or open area in the chassis front panel (Figure 26) or by an airflow duct (Figure 27).

Figure 26: Thermal Module Interface Provided by Vented Front Chassis Panel



OM16730

Figure 27: Thermal Module Interface Provided by a Front Panel or Side Panel Airflow Duct



OM16731

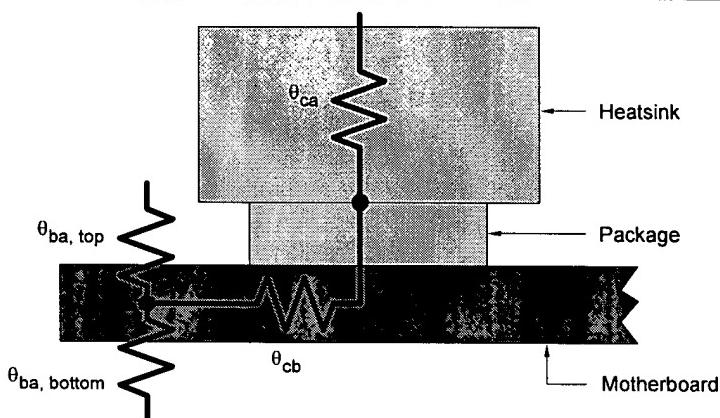
The *BTX Interface Specification* also requires that the Thermal Module Interface be sealed in a way that prevents recirculated air from entering it. This assures that the Thermal Module fan will only pull air from outside the chassis through the Thermal Module Interface. A Thermal Module Interface inside the chassis walls (that is, if it is not created by the chassis sheet metal front panel) must be provided by the end of an airflow duct. This airflow duct should start at a chassis panel (Figure 27 illustrates start positions at both the front panel and a side panel) and end at the Thermal Module Interface. It would, of course, be beneficial to design gradual turns into any duct design, as opposed to the sharp turn shown in the rightmost illustration in Figure 27.

Should the chassis sheet metal front panel provide the Thermal Module Interface, the chassis supplier should appropriately engineer the vent open area. The open area is typically constrained by electromagnetic compatibility (EMC) that the chassis must also provide. Recall that the CPU is located near the front of the motherboard and, therefore, near the front panel and that the proximity of this electromagnetic interference (EMI) source to an open area in the front panel may be a concern. Nonetheless, the chassis supplier should optimize the hole size, shape, and pattern to create effective EMC without introducing airflow impedance that will substantially degrade the Thermal Module ψ_{ca} or operating point airflow. In this design situation, it is important for the system integrator to understand the impact that the near field impedance of the front panel vent will have on these Thermal Module characteristics. In particular, it may be beneficial to offset the vent area from the front panel to increase the distance between the vent and the Thermal Module fan. A Thermal Module Interface is required by the *BTX Interface Specification*. The opening and physical interface requirements described must be provided by the BTX chassis supplier. It is against the Thermal Module Interface physical features that the inlet of the Thermal Module will fit. The system integrator will determine the quality of the seal between the Thermal Module and Thermal Module Interface, although the Thermal Module supplier will typically provide a compliant physical interface (for instance, a compression gasket or seal) to accommodate the expected dimensional tolerance variation.

2.4.6 MCH Airflow and Heat Transfer

Figure 28 illustrates the heat transfer path and associated network thermal resistance diagram for an MCH. Intel typically offers the MCH in a Flip Chip Ball Grid Array (FCBGA) package that does not have a metal lid, molded plastic, or other substance over the bare silicon die. Silicon is not a particularly good heat conductor (from Figure 28, θ_{ca} would be very high in the absence of an MCH heatsink) but the die surface is available for heat transfer augmentation, so a heatsink is typically directly attached to its surface. The heatsink significantly reduces the resistance of this heat transfer path and heat flow from the top surface of the MCH package is the primary heat transfer path.

Figure 28: MCH Heat Transfer Path and Resistance Network Diagram



OM16732

As discussed in Section 2.4.3, the below motherboard airflow augments the heat transfer for the MCH. It is also true that the above motherboard airflow flows over the motherboard area near the MCH and this improves the convective heat transfer from the motherboard topside, as well. Both of these motherboard convective heat transfer paths augment the heat transfer from the MCH and lower its temperature. However, because the resistance of the heat transfer path through the MCH heatsink is so low, the relative importance of heat transfer into and out of the motherboard diminishes.

Nonetheless, the heat transfer capability of the motherboard in the area of the MCH can and should be characterized. Knowledge of these heat transfer paths can be used to optimize the performance and cost of the MCH heatsink. The network resistance diagram in Figure 28 can be used to define the total heat transfer resistance equation (Equation 9) – in this equation, the “up” and “down” terms are in reference to the power source (MCH silicon). The required total resistance can be calculated from the component power and temperature requirements, and the MCH ambient temperature (Equation 10).

Note: MCH power (Q_{MCH}) and case temperature (T_{case}) requirement are specified in the component Datasheet. Ambient temperature ($T_{ambient}$) in the MCH region should be extracted from a system level thermal numerical model or from system temperature measurements. Ambient temperatures characterized in Intel reference design systems are available in Section 2.7.3.

Equation 9: Total Heat Transfer

$$\begin{aligned}\theta_{\text{total}} &= 1/(1/\theta_{\text{up}} + 1/\theta_{\text{down}}) \\ &= 1/((1/\theta_{\text{ca}}) + 1/(\theta_{\text{cb}} + 1/(1/\theta_{\text{Ba, top}} + 1/\theta_{\text{ba, bottom}})))\end{aligned}$$

Equation 10: Required Total Resistance

$$\theta_{\text{total}} = (T_{\text{case}} - T_{\text{ambient}}) / Q_{\text{MCH}}$$

From this equation, the required heatsink resistance (θ_{ca}) can be calculated (Equation 11).

Equation 11: Total Heatsink Resistance

$$\theta_{\text{ca}} = 1/((Q_{\text{MCH}}/(T_{\text{case}} - T_{\text{ambient}}) - 1/(\theta_{\text{cb}} - 1/(1/\theta_{\text{Ba, top}} + 1/\theta_{\text{ba, bottom}})))$$

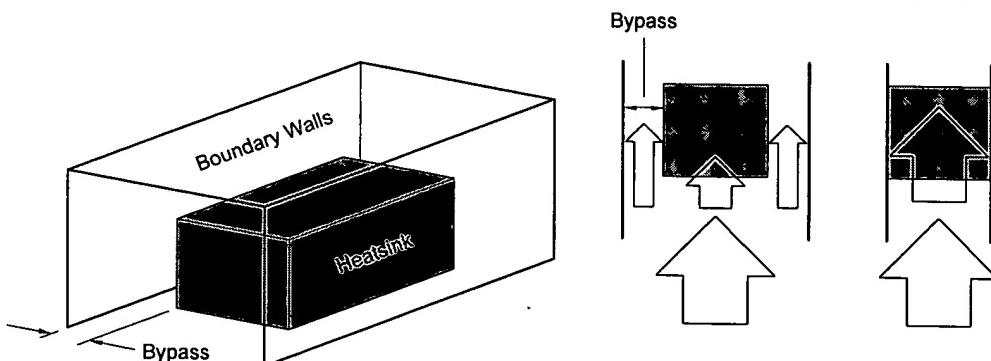
Equation 11 shows how the required MCH heatsink performance must improve if the motherboard heat transfer characteristics are unknown. Table 2 offers an example of the relative contribution of the θ_{ba} resistance terms – when the motherboard resistance terms are unknown, they must be assumed to be infinite resistance.

Table 2: Motherboard Heat Transfer Impact to Required MCH ψ_{ja}

Case Temperature (°C)	112		
Ambient Temperature (°C)	45		
MCH Power (W)	18		
θ_{cb} (°C/W)	$\theta_{\text{ba, top}}$ (°C/W)	$\theta_{\text{ba, bottom}}$ (°C/W)	θ_{ca} (°C/W)
15	15	15	3.59
30	30	40	3.66
∞	∞	∞	3.72

The performance and cost for the MCH heatsink is also a function of the airflow through the MCH heatsink. The Thermal Module exhaust airflow is the primary inlet airflow for the MCH heatsink. As discussed in Section 2.4.1, the heat transfer performance of a heatsink is a strong function of the amount of airflow that passes through it. However, unless it is directed through a heatsink, airflow will often choose to go around it (see Figure 29).

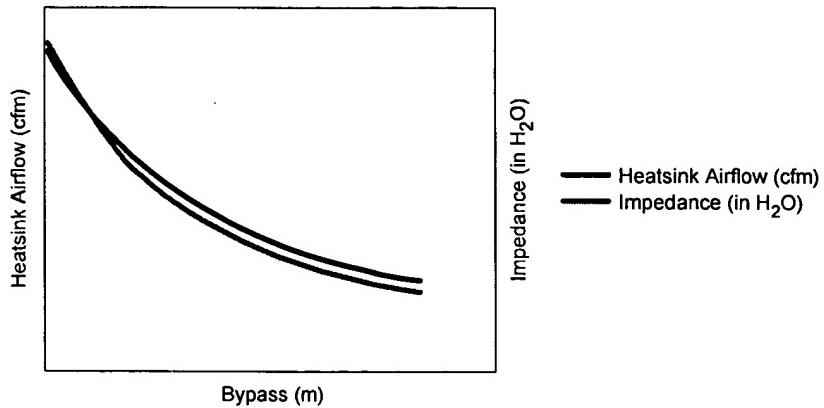
Figure 29: Illustration of Airflow through a Heatsink as a Function of Bypass



OM16733

The amount of open channel area around a heatsink (termed bypass) determines the amount of airflow that will go through a heatsink. As illustrated in Figure 30, reducing the bypass around a heatsink increases the amount of airflow that will go through it. Although the airflow through a heatsink improves with a reduction in bypass, the impedance of that restricted airflow path also increases (Figure 30).

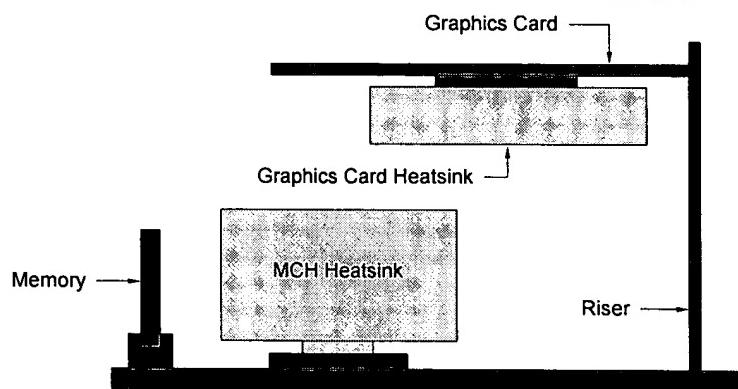
Figure 30: Impact of Bypass on Heatsink Airflow and Impedance



OM16734

A system designer will need to be cognizant of the bypass around the MCH in the design or selection of the MCH heatsink, since there are several system configuration choices that will impact the airflow bypass around the MCH. For instance, an add-in card on a riser may occupy the volume immediately above the MCH (above the Zone B Keepout height required in the *BTX Interface Specification*), as shown in Figure 31 or the area above the MCH may be free all the way to the chassis boundary wall, as shown in Figure 32.

Figure 31: MCH Bypass Illustration – Graphics Add-in Card on a Riser



OM16735